



(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 938 703 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
02.07.2003 Bulletin 2003/27

(21) Application number: **97951436.1**

(22) Date of filing: **13.11.1997**

(51) Int Cl.⁷: **G06F 11/00, G06F 9/38**

(86) International application number:
PCT/US97/20980

(87) International publication number:
WO 98/021655 (22.05.1998 Gazette 1998/20)

(54) REAL TIME PROGRAM LANGUAGE ACCELERATOR

ECHTZEITPROGRAMM-SPRACHBESCHLEUNIGER

ACCELERATEUR DE LANGAGE DE PROGRAMMATION TEMPS REEL

(84) Designated Contracting States:
**AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE**

(30) Priority: **13.11.1996 US 30688 P**

(43) Date of publication of application:
01.09.1999 Bulletin 1999/35

(73) Proprietor: **Paran, Arik
Sunnyvale, CA 94087 (US)**

(72) Inventor: **RAZ, Yair
Sunnyvale, CA 94087 (US)**

(74) Representative: **Steil, Christian, Dipl.-Ing. et al
Witte, Weller & Partner,
Rotebühlstrasse 121
70178 Stuttgart (DE)**

(56) References cited:
**EP-A- 0 154 529 US-A- 4 205 370
US-A- 4 674 089 US-A- 4 980 821
US-A- 5 126 541 US-A- 5 420 989
US-A- 5 442 777**

- SHAND M ET AL: "HARDWARE SPEEDUPS IN LONG INTEGER MULTIPLICATION" COMPUTER ARCHITECTURE NEWS, US, ASSOCIATION FOR COMPUTING MACHINERY, NEW YORK, vol. 19, no. 1, 1 March 1991 (1991-03-01), pages 106-113, XP000201929 ISSN: 0163-5964
- SILBERMAN G M ET AL: "AN ARCHITECTURAL FRAMEWORK FOR MIGRATION FROM CISC TO HIGHER PERFORMANCE PLATFORMS" INTERNATIONAL CONFERENCE ON SUPERCOMPUTING, CONFERENCE PROCEEDINGS, 19 July 1992 (1992-07-19), XP000576925
- MICROPROCESSOR REPORT, March 1996, Vol. 10, No. 4, CASE B., "Implementing the Java Virtual Machine; Java's Complex Instruction Set Can Be Built in Software or Hardware", pages 12-18, XP000987276.
- BYTE, November 1996, Vol. 21, No. 11, WAYNER P., "Sun Gambles on Java Chips", pages 79-85, XP000641070.

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description**TECHNICAL FIELD**

5 [0001] The present invention relates to the field of computer processing and more specifically to the real time interpretation and operation of computer code using a combination of unique hardware and software. The predominant current usage of the present inventive real time program language accelerator is in the execution of Java™ code wherein it is desirable to be able to run such code at high execution speed on a variety of different processors.

BACKGROUND ART

10 [0002] It is known in the art to provide interpreters for converting higher level computer languages into machine readable code in near real time. However, such interpreters must necessarily slow down the operation of the program being run at least at some point in during the loading and running of the program. In particular, regarding the Java™ program language. the several available solutions for running Java™ programs, including the " Java™ Virtual Machine" 15 are software based programs which require execution in order to accomplish their tasks.

15 [0003] It would be beneficial to have a method and/or means for running Java™ code on any available type of processor. Also, it would be beneficial to have some method and/or means for processing the Java™ code such that it will actually run faster than it otherwise would. or at least that would not slow down the overall process of interpreting 20 and running the Java™ code. However, to the inventor's knowledge, no system for accomplishing these objectives has existed in the prior art. Indeed, although software solutions such as the Java™ Virtual machine do attempt to optimize the operation of the code, a limitation has been the characteristics of the particular processor and system on which the code is to be run.

DISCLOSURE OF INVENTION

25 [0004] Accordingly, it is an object of the present invention to provide a method and means for easily running a non-native program language on virtually any type of processor.

30 [0005] It is still another object of the present invention to provide a method and means for accelerating the operation of Java™ program code.

[0006] It is yet another object of the present invention to provide a method and means for translating and executing Java™ code which is easily implemented in hardware.

[0007] It is still another object of the present invention to provide a method and means for translating and executing Java™ code which will not detract from the ability of the computer system to execute non- Java™ code.

35 [0008] Briefly, the preferred embodiment of the present invention is an integrated circuit "program language accelerator core" which can be embodied as part of a CPU chip, on a separate chip. or even on a separate board. The program language accelerator core has a dedicated hardware stack memory for directly providing the stack memory which is required to be either present or emulated for the execution of Java™ code. A direct memory access ("DMA") controller is also provided for shifting data into and out of the stack memory as overflow and underflow conditions (or 40 other specified conditions) occur. A software portion of the invention translates the Java™ code into a machine native language code. and also optimizes the code to perform unique inventive functions as required. such as writing to a memory address which is predetermined to perform specified functions. In this manner, operations that might otherwise take several clock cycles can be performed in a single (or at least fewer) clock cycles.

[0009] An advantage of the present invention is that the speed of execution of Java™ code is greatly increased.

45 [0010] A further advantage of the present invention is that Java™ code can be readily executed on essentially any type of processor.

[0011] Yet another advantage of the present invention is that it can be easily and inexpensively implemented such that even affordably priced computers can be optimized for the execution of Java™ code.

50 [0012] Still another advantage of the present invention is that it is not difficult or expensive to adapt to new types of processors such as might become available.

[0013] These and other objects and advantages of the present invention will become clear to those skilled in the art in view of the description of the best presently known mode of carrying out the invention and the industrial applicability of the preferred embodiment as described herein and as illustrated in the several figures of the drawing.

BRIEF DESCRIPTION OF THE DRAWINGS**[0014]**

- 5 Fig. 1 is a block diagram of an integrated circuit having thereon a program language accelerator core according to the present invention;
- Fig. 2 is a block diagram showing an example of how an CPU enhanced according to the present invention might be integrated into a computer system;
- 10 Fig. 3 is a more detailed block diagram of the intelligent stack of Fig. 1;
- Fig. 4 is a memory map of the intelligent stack depicted in Figs. 1 and 3;
- Fig. 5 is a flow diagram of one aspect of the inventive process;
- Fig. 6 is a flow diagram of a value push operation according to the present invention;
- Fig. 7 is a flow diagram of a value pop operation according to the present invention;
- 15 Fig. 8 is a flow diagram of a constant push operation according to the present invention;
- Fig. 9 is a flow diagram of an arithmetic operation according to the present invention;
- Fig. 10 is a flow diagram of a local variables load operation according to the present invention; and
- Fig. 11 is a flow diagram of a local variables store operation according to the present invention.

BEST MODE FOR CARRYING OUT INVENTION

- 20 **[0015]** The best presently known mode for carrying out the invention is a real time accelerator core. The inventive real time accelerator core is depicted in a block diagrammatic view in Fig. 1 and is designated therein by the general reference character 10.
- 25 **[0016]** The real time accelerator core 10 has an intelligent stack 12 and a smart DMA controller 14. The real time accelerator core 10 will allow real time translation and execution of Java™ object code by any target native CPU 16 such that the native CPU can execute Java™ programs without any additional software based translation or interpretation in a much higher performance mode (compared to the standard Java™ interpreter or JIT compilers). It should be noted that, in the example of Figure 1, the program language accelerator core 10 and the CPU are embodied on a single CPU chip 18 although, as previously discussed herein, this is not a necessary aspect of the invention.
- 30 **[0017]** Figure 2 is a diagrammatic example of a typical computer configured for operation with the CPU chip 18 having therein the program language accelerator 10. It will be noted that the connections of the CPU chip 18 within the computer are not significantly different from those of a comparable conventional prior computer (not shown). A data bus 22, an address bus 24 and a control bus 26 are each provided with the appropriate data paths 28 for communicating with the CPU chip 18, a (RAM) memory 30 and an I/O section 32.
- 35 **[0018]** The data bus 22 is the native CPU data bus. The width of the data bus 22 should be the natural data bus width of the native CPU 16 (8, 16, or 32 bits). The interface logic in the program language accelerator core 10 will take care of any data width related tasks. The address bus 24 is the native CPU address bus. The width of the address bus 28 should be the natural address bus width of the native CPU 16. The control bus 26 will carry the several control signals that can be found in any CPU: Clock, Reset, Read and Write, interrupt lines, bus request, etc.
- 40 **[0019]** It should be noted that the example of Figure 2 is provided only to show the context in which inventive program language accelerator core 10 might be used, and is not intended to disclose an inventive aspect of the invention. In operation, the memory 30 will contain the required initialization software for the native CPU 16 (Figure 1), an operating system (if used), I/O device drivers or control code, as well as a program language accelerator software 34 the functions of which will be discussed in greater detail hereinafter. In addition, the memory 30 can contain any compiled Java™ code required (specific custom class libraries, for example).
- 45 **[0020]** Depending on the end user's application, the computer 20 equipped with the inventive program language accelerator core 10 and the program language accelerator software 34 can execute any Java™ code which already exists in non-volatile memory, or load it through the applicable I/O devices (communication ports, disk, etc.).
- 50 **[0021]** The program language accelerator core 10 takes advantage of the fact that the Java™ Virtual Machine is a stack based machine emulation, and most stack operations on a contemporary CPU (such as the native CPU core 16 of Figure 1) take more than 1 clock cycle to complete. The intelligent stack 12 is essentially a large virtual cache for the Java™ Virtual Machine stack. The DMA controller 14 (Figure 1) keeps the cache consistent if a stack overflow or underflow would occur due to the limited size of the actual cache, as will be discussed in greater detail hereinafter.
- 55 **[0022]** Figure 3 is a more detailed block diagram of the intelligent stack 12 previously discussed herein in relation to Figure 1. In the view of Figure 3 it can be seen that the intelligent stack 12 has a simple internal arithmetic logic unit ("ALU") which allows the pushing of constants as defined by the Java™ Virtual Machine specification, along with the operations that update the tops of the stack such as *iadd* or *iinc*, all of which will be familiar to one skilled in the implementation of the Java™ Virtual Machine. A cache 38 is the actual cache memory within the intelligent stack. A

stack controller 40 has a stack pointer 42 and stack pointer next 44 such as are found in conventional stack control devices. A conventional decoder 46 is provided for decoding addresses and instructions and an output MUX 48 provides output from the source within the intelligent stack 12 which is appropriate to the operation.

[0023] The intelligent stack 12 is a 64 x 32-bit register file configured to be used as a stack. The intelligent stack 12 is memory mapped into the native CPU 16 memory address space, and allows the native CPU 16 fast stack operation (push. pop), and also random access to any one of the registers. The intelligent stack 12 provides circuitry for automatic stack overflow/underflow detection. The smart DMA controller 14 is configured to read and write the content of the entire intelligent stack 12 to or from the memory 30, such as when it is required to shift between different tasks. The smart DMA controller 14 also corrects potential stack overflow or underflow by temporarily storing any excess of data from the intelligent stack 12 to the memory 30. The smart DMA will dump a block of words to main memory 30 when the stack approaches overflow (on Push operation), or will load a block of words from main memory 30 when the stack approaches underflow (on Pop operation). The smart DMA controller can load a local variable from main memory 30 (on Variable Load miss), store a local variable to main memory 30 (on Variable Store miss), dump the entire cache 38 to main memory 30 to prepare for a context switch. or load the entire cache 38 from memory 30 to perform a context switch. The smart DMA can also optionally be used to accelerate thread context switching by moving data into and out of the cache 38, as required.

[0024] It should be noted that the inventive program language accelerator 10 could be made to operate without the smart DMA controller 14, although the addition of the smart DMA controller 14 does significantly increase the usefulness and functionality of the program language accelerator 10.

[0025] Since Java™ operations are based on a stack machine architecture, the program language accelerator software 34 is required to perform only a very simple translation of the Java™ code to native CPU code. The translation is done using the support language accelerator software 34, preferably as the Java™ code is loaded. Optimization according to the present invention will be provided during the translation by methods including the present inventive method of providing that certain operations be directed to certain memory addresses, as will be discussed in more detail hereinafter. Since the native CPU 16 does not need to deal with stack management (which is handled by the intelligent stack 12), very high speed execution of Java™ programs is possible.

[0026] According to the present invention data will come from the data bus 22 and flow onto the top of the stack (much as a conventional stack operation) or else will have an arithmetic operation including the value at the top of the stack and then replace or push onto the top of the stack. Unique to the present invention is the fact that all of the operations are determined by the addresses used. The intelligent stack 12 occupies four times the address space of the cache 38. That is, four times the number of addresses are mapped to memory of the CPU chip 18 than are required to fully address all of the actual memory locations in the intelligent stack 12. The size of the cache in the best presently known embodiment 10 of the present invention is 64 words so the intelligent cache 12 occupies 256 locations in this embodiment.

[0027] Figure 4 is a simple memory map 50 of the memory addresses occupied by the intelligent stack 12. In the view of Figure 4 it can be seen that the memory map 50 has a first area 52, a second area 54, a third area 56 and a fourth area 58. Of these, only the first area 52 relates directly to the physical cache 38. A complete address 60 for addressing the memory map 50 will have two select area bits 62 for selecting which of the areas 52, 54, 56 or 58 is to be addressed and an 8 bit select byte 64 for addressing specific data locations within the respective areas. Since each of the four areas 52 allows different operations to be executed by reading or writing to an address therewithin, the upper two address bits (the select area bits 62) of program language accelerator core 10 determines which area 52, 54, 56 or 58 is addressed, and so determines the major operation mode of the device.

[0028] Table 1, below, provides a general description of the operations which are performed when the native CPU 16 addresses each of the four areas 52, 54, 56 and 58 of the memory map 50.

TABLE 1

Upper Address Bits	program language accelerator core Major Operation Areas
00	Direct access to hardware stack (cache) data.
01	Access to Configuration and Control Registers
10	Read: Pop value from stack. Write: Push value or constant into stack (depending upon address within the area).
11	Read: Undefined. Write: Replaces TOS with the required operation's result. Operation is performed on TOS and the data written. Operation depends upon address within the area.

[0029] Briefly, values written to the first area 52 of the memory map 50 may be directly written or read by the native CPU 16 at the base address of the device. The second area 54 is used for the configuration and control registers. The third area 56 is used for push and pop operations. The fourth area 58 is used for arithmetic operations that replace the value at the top of the stack. The decoder 46 looks at the incoming address request and thereby automatically determines the operation, the correlation between address and operation having been previously established during translation by the program language accelerator software 34. The data busses and timing are generic single cycle busses where the address, read and write signals are valid on the clock, and read data is expected by the following clock.

[0030] Figure 5 is a flow diagram 68 depicting the general operation of the computer 20 according to the present inventive method. Whenever the computer 20 is restarted, specific initialization is done by the native CPU 16: Setting up hardware devices, initializing the operating system (if used), executing any required startup program, and so on. These programs are written in the native CPU 16 language (e.g. instruction set) and are executed by the native CPU 16 directly. These are done without any connection to Java™ and/or program language accelerator core 10. Also, in a set up internal registers operation 70 the native CPU will initialize memory registers, including those that are set aside for addressing the program language accelerator core 10. This is the initial phase of operation when the user invokes and starts the Java™ Execution Environment. During initialization of the Java™ Virtual Machine, the Java™ Virtual Machine will detect the existence of program language accelerator core 10 and its support software (the program language accelerator software 34), and initializes both of them. Among the various initializations done at this phase, some of the most important are setting the various Configuration Registers of program language accelerator core 10 (as described in herein in relation to the second area 54 of the memory map 50).

[0031] Any required native program might then be executed by the computer 20. This is also done directly by the native CPU 16 without any operation required by program language accelerator core 10. (These operations are not shown in the flow diagram 68 of Figure 5, and are mentioned here only to put the flow diagram 68 in the correct context.)

[0032] When a Java™ class file needs to be executed (either it exists already in memory, or it needs to be loaded from disk or received from a communication line), the Native CPU 16 uses the JVM class loader and the program language accelerator software 34 to load, prepare, perform translation, and start executing the required file. This process is made of several steps which will be described in details in the following sections: In a class file load operation 72, a Java™ class file is loaded. This part is executed entirely by the standard Java™ Virtual Machine code, as described in "The Java™ Virtual Machine Specification". In a linking and verification operation 74, linking, verification, preparation and resolution are performed conventionally by the Java™ Virtual Machine linker. This part is also executed entirely by the standard Java™ Virtual Machine code.

[0033] Following, in a translation operation 76, the native CPU 16 locates the class file's implementation code (e.g. the Java™ Virtual Machine byte code which implements the class), translates it to native instructions, and loads it to an execution area in the memory 30. This part is done entirely by the program language accelerator software 34. The purpose of the translation phase is to convert the Java™ byte code (which was loaded and linked in the previous phases) to the native CPU instruction streams, suitable to operate program language accelerator core 10. Since, as described herein, reading or writing from/to special memory locations invokes all of program language accelerator core 10 operations, the translated native code will mostly contain read and write instructions from/to the various memory areas of program language accelerator core 10. The specific translation of byte code to native code depends on the native CPU that is attached to program language accelerator core 10. The translated code is stored in the memory 30.

Once a class file is fully translated, the original byte code image can be discarded, and only the translated native code will be used.

[0034] Once the loading process is complete, the native CPU 16 will perform a branch (jump) to the entry point of the translated native code and will start executing the class initialization code, which is now a native program with instruction sequences that take advantage of program language accelerator core 10 dedicated hardware stack and logic, as discussed throughout this disclosure. This is indicated by an execute native code operation 78 in the flow diagram 68 of Figure 5.

[0035] Further details of the operation of the inventive method occurring within the execute native code operation 78 will be discussed in relation to additional flow diagrams hereinafter. These operations and respective flow diagrams are as follows: A value push operation 80 (Figure 6) and a value pop operation 82 (Figure 7), both of which relate to operations which occur when the native CPU 16 addresses the second area 54 of the memory map 50 of the intelligent stack 12. A constant push operation 84 (Figure 8) also relates to operations which occur when the native CPU addresses selected locations of the second area 54 of the memory map 50 of the intelligent stack 12. (One familiar with Java™ will recognize that there is no need for an equivalent "constant pop" operation.)

[0036] An arithmetic operations 86 flow diagram (Figure 9) describes operations which occur in the program language accelerator core 10 when selected areas of the third area 54 of the memory map 50 are addressed. A local variables load operation 88 (Figure 10) and a local variables store operation 90 (Figure 11) describe these respective functions and will provide more detail about the operation of the smart DMA controller 14.

[0037] Returning now again to a discussion of the different functioning of the four areas 52, 54, 56 and 58 of the

memory map 50 introduced previously herein, it will be remembered that, in the best presently known embodiment 10 of the present invention, the program language accelerator core 10 contains 64 words in its hardware intelligent stack 12. This address space is divided into the four major areas 52, 54, 56 and 58, each of which allows different operations to be executed by reading or writing to an address therewithin. Regarding the first area 52 of the memory map 50, this area acts as one skilled in the art would expect the memory map 50 to behave were this the only area of the memory map - that is, if there were a one to one correlation between the addresses of the memory map 50 and the cache 38. The first area 52 of the memory map 50 is provided to allow read and write access to any of the hardware stack registers (64 32-bit registers) as though they were random access memory.

[0038] Regarding the second area 54 of the memory map 50, the registers in this area control the operation of the program language accelerator core 10. These registers have read and write access. This area contains four special write-only locations, also. Two of the write only locations are used for variable operations and two are used for context switching. Addressing any of these registers is done via the low address bits of the native CPU 16. Table 2, below is a listing of the applicable registers within the second area 54 of the memory map 50.

TABLE 2

Address	Register	Register or Function
0	SP	Stack Pointer.
1	BOSP	Bottom of Stack Pointer.
2	WLIMIT	Upper limit in main memory where the stack can be written to.
3	RLIMIT	Lower limit in memory where the stack can be read.
4	VREG	Variable Address Register. This location can be written or read for testing purposes. It is automatically written with the address of the last accessed variable.
5	SWAPINSIZE	Swap-in Size. This register controls how many words will be read in when a swap-in DMA cycle is started.
8	VLOAD	Variable Load. This register has write only access. When this special location is written, the data bus has the absolute address of a local variable that should be pushed on to the top of the stack. A DMA cycle is started if the variable is not in the cache. If the variable is in the cache, one additional cycle is taken to transfer the contents to the top of the stack
9	VSTOTE	Variable Store. This register has write only access. When this special location is written, the data bus has the absolute address of local variable that should get the value from the top of the stack. A DMA cycle is started if the variable is not in the cache. If the variable is in the cache, one additional cycle is taken to transfer the contents to the top of the stack
14	SWAPIN	Swap-in. This register has write only access. When this special location is written, the value on the data bus is ignored and a block mode DMA cycle is started to fill the cache from main memory. The number of word read depends on the value of the SWAPINSIZE register. so this may be tuned for a specific state of a thread
15	SWAPOUT	Swap-out. This register has write only access. When this special location is written, the value on the data bus is ignored and a block mode DMA cycle is started to flush the cache to main memory. The number of word read depends on the value of SP and BOSP

[0039] Regarding the third area 56 of the memory map, this area is used for pushing values or constants to the stack or popping values from the stack. The operation to be performed is determined by the low address bits of the processor (the select byte 62). That is, the address determines the operation to be performed. These functions are listed below in table 3. One skilled in the art will recognize that, in some instances, the value of data provided by the native CPU 16 to the program language accelerator core 10 will be irrelevant, since the operation is performed using a constant value.

TABLE 3

Address	Operation	Operation Description
0	PUSHPOP	Pushes the value on the data bus to the stack (write operation) or pops a value from TOS to the data bus (read operation).
2	ICONSTM1	Write only location. Pushes a constant integer (-1)
3	ICONST0	Write only location. Pushes a constant integer 0
4	ICONST1	Write only location. Pushes a constant integer 1
5	ICONST2	Write only location. Pushes a constant integer 2
6	ICONST3	Write only location. Pushes a constant integer 3
7	ICONST4	Write only location. Pushes a constant integer 4
8	ICONST5	Write only location. Pushes a constant integer 5
11	FCONST0	Write only location. Pushes a constant float 0.0
12	FCONST1	Write only location. Pushes a constant float 1.0
20	FCONST2	Write only location. Pushes a constant float 2.0

[0040] Regarding the fourth area 58 of the memory map 50, this area is used to initiate arithmetic operations on the value at the Top-Of-Stack. The value at the top of the stack is replaced with the result of an arithmetic operation between value on the data bus the current value at the Top-Of-Stack. The arithmetic operation to be performed is determined by the low address bits of the processor (the select byte 62). All of these operations are performed on integer values (32-bit integers). These functions are listed in table 4, below.

TABLE 4

Address	Operation	Arithmetic Operation Description
16	IADD	Write only location. Adds the value at the top of the stack to the value on the data bus, replacing the value at the top of the stack.
17	ISUB	Write only location. Subtracts the value on the data bus from the value at the top of the stack, replacing the value at the top of the stack.
18	INEG	Write only location. Subtracts the value at the top of the stack from 0, replacing the value at the top of the stack. Ignores the value on the data bus.
19	IOR	Write only location. Performs a bit-wise OR on the value at the top of the stack with the value on the data bus, replacing the value at the top of the stack.
20	IAND	Write only location. Performs a bit-wise AND on the value at the top of the stack with the value on the data bus, replacing the value at the top of the stack.
21	IEXOR	Write only location. Performs a bit-wise Exclusive-OR on the value at the top of the stack to the value on the data bus, replacing the value at the top of the stack
22	IINCR	Write only location. Adds 1 to the value at the top of the stack, replacing the value at the top of the stack. Ignores the value on the data bus.

[0041] Referring now to the value push operation 80 shown in Figure 6, when the program language accelerator software 34 has translated and optimized the Java™ code regarding a value push operation, and it is time for such operation to be executed, the value will be written to the location "pushpop" (see Table 3) in operation 80a. In an operation 80b, the address is decoded (by the decoder 46 - Figure 3) the value is so directed. Then, in an operation 80c the value written from the data bus (Figure 2) is written into the hardware stack (the cache 38 - Figure 3) at the location pointed by its TOS register, and the TOS is incremented. If the hardware stack (cache 38) approaches overflow, as determined in a decision operation 80d, then the smart DMA controller 14 is initiated to save a portion of the content of the cache 38 to the memory 30 (Figure 2) in an operation 80e.

[0042] Referring now to the value pop operation 82 shown in Figure 7, when the program language accelerator

software 34 has translated and optimized the Java™ code regarding a value pop operation, and it is time for such operation to be executed, the read command will be directed to the appropriate location in operation 82a. In an operation 82b, the address is decoded (by the decoder 46 - Figure 3). Then, in an operation 82c the value read (popped) from the cache 38 (Figure 3) and the TOS pointer is decremented. The value is sent to the data bus (Figure 2) in an operation 82d. If this leaves the hardware stack (cache 38) in an underflow condition, as determined in a decision operation 82e (that is, if the cache 38 has reached a predetermined level of non-use), then the smart DMA controller 14 is initiated to restore a portion of the content of the cache 38 from the memory 30 (Figure 2) in an operation 82f.

[0043] Referring now to the constant push operation 84 shown in Figure 8, when the program language accelerator software 34 has translated and optimized the Java™ code regarding a constant push operation, and it is time for such operation to be executed, the write command will be directed to the appropriate location in operation 84a. In an operation 84b, the address is decoded (by the decoder 46 - Figure 3). Then, in an operation 84c the constant value dictated by the particular address selected is written (pushed) into the hardware stack (the cache 38 - Figure 3) at the location pointed by its TOS register, and the TOS is incremented. If the hardware stack (cache 38) approaches overflow, as determined in a decision operation 84d, then the smart DMA controller 14 is initiated to save a portion of the content of the cache 38 to the memory 30 (Figure 2) in an operation 84e.

[0044] Referring now to the arithmetic operation 86 shown in Figure 9, when the program language accelerator software 34 has translated and optimized the Java™ code regarding a particular arithmetic operation, and it is time for such operation to be executed, the write command will be directed to the appropriate location in operation 86a. Note that exactly what that appropriate location might be will be dictated by the particular arithmetic operation that is to be performed, as listed previously herein in Table 4. In an operation 86b, the address is decoded (by the decoder 46 - Figure 3) and the corresponding data is fetched from that data bus 22 (Figure 2). Then, in an operation 86c the arithmetic operation corresponding to the selected address is performed on the value at the top of slack using the value written from the data bus (Figure 2) and the result is written into the hardware stack (the cache 38 - Figure 3) at the location pointed by its TOS register.

[0045] In a variable store operation 90 (Figure 11), when the program language accelerator software 34 has translated and optimized the Java™ code regarding a variable store operation, and it is time for such operation to be executed, a write will be directed to the dedicated memory location corresponding to the store operation in the intelligent stack 12 in operation 90a. The value written to the data bus 22 should be the absolute memory address of the required local variable to be stored. In an operation 90b, the address is decoded (by the decoder 46 - Figure 3). Then, if the required variable is in the stack (as determined in a decision operation 90c) the variable value is read from the TOS and stored to its required address and the top of stack is decremented. If the required variable is not in the stack, as determined in the decision operation 88c, then the smart DMA controller 14 is initiated to store the variable value from the top of stack and decrement the TOS in an operation 90e. If the hardware stack (cache 38) approaches underflow, as determined in a decision operation 90f, then the smart DMA controller 14 is initiated to restore the stack content from the main memory 30.

[0046] In a variable load operation 88 (Figure 10), when the program language accelerator software 34 has translated and optimized the Java™ code regarding a variable load operation, and it is time for such operation to be executed, a write will be directed to the dedicated memory location in the intelligent stack 12 in operation 88a which is dedicated to the variable store operation. The value written to the data bus 22 should be the absolute memory address of the required local variable to be stored. In an operation 88b, the address is decoded (by the decoder 46 - Figure 3). Then, if the required variable is in the stack (as determined in a decision operation 88c) the variable value is read from the cache 38 and placed at the top of stack in an operation 88d - and the TOS is incremented. If the required variable is not in the stack, as determined in the decision operation 88c, then the smart DMA controller 14 is initiated to load the variable value to the top of stack and increment the TOS in an operation 88e. If the hardware stack (cache 38) approaches overflow, as determined in a decision operation 88f, then the smart DMA controller 14 is initiated to transfer the stack content to the main memory 30 from the hardware stack (cache 38).

[0047] Various modifications may be made to the invention without altering its scope. For example, although the inventive program language accelerator core 10 and associated program language accelerator software 34 are described herein as being optimized for use with the Java™ program language, the principles involved are equally applicable for use with other program languages, particularly if such languages might be developed primarily for use with stack based systems.

[0048] As previously mentioned, yet another likely modification would be to implement the program language accelerator core as a device physically distinct form the CPU chip 18 such that it could more readily be added to existing systems or existing system designs.

55

INDUSTRIAL APPLICABILITY

[0049] The inventive program language accelerator core 10 and the related program language accelerator software

34 are intended to be widely used for the real time execution of Java™ code in conjunction with processors which are otherwise optimized for executing programs written in languages other than Java™. As can be appreciated in light of the above description, the program language accelerator core 10 uses an entirely new concept in that the program language accelerator software, when interpreting the Java™ code for execution in the native code of the native CPU 16 will cause certain operations to be directed to certain virtual memory addresses such that the operation is automatically accomplished by the intelligent stack 12. That is, the intelligent stack 12 will know what operation is to be performed and will perform it based solely on the address to which the data is written in the memory map. This will save from one to several clock cycles per operation where this feature is invoked. Thereby, the speed of execution of Java™ code will be greatly enhanced without burdening the computer 20 with the running of a virtual machine, or the like, in the background.

[0050] Since the program language accelerator core 10 of the present invention may be readily produced and integrated into existing designs for systems and devices, and since the advantages as described herein are provided, it is expected that it will be readily accepted in the industry. For these and other reasons, it is expected that the utility and industrial applicability of the invention will be both significant in scope and long lasting in duration.

[0051] All of the above are only some of the examples of available embodiments of the present invention. Those skilled in the art will readily observe that numerous other modifications and alterations may be made without departing from the scope of the invention. Accordingly, the above disclosure is not intended as limiting and the appended claims define the scope of the invention.

Claims

1. A computer system providing accelerated processing of stack oriented interpretive language instructions, said system (20) comprising:

a native processing unit (16) for executing native language instructions;
a translator (34) for translating the interpretive language instructions to native language instructions,

characterized by said translator (34) being capable of generating selected address values (60, 62, 64) associated with selected ones of said native language instructions, each of said selected instructions and said corresponding selected address values (60, 62, 64) being associated with a corresponding processing accelerator core operation;

a processing accelerator core (10) coupled for communication with said native processing unit (16), said processing accelerator core (10) including,

a stack cache memory device (38) having a plurality of memory address locations indicated by corresponding ones of said selected address values (60, 62, 64), said memory address locations for storing data associated with corresponding ones of said processing accelerator core operations,

core control logic (14, 40, 46) responsive to said selected address values (60, 62, 64), and being operative to control writing and reading of said data associated with said selected native language instructions to and from said stack cache memory device (38), and

an arithmetic logic unit (36) for receiving said data from corresponding ones of said memory address locations of said stack cache memory device (38), and being operative to perform said processing accelerator core operations using said data, each of said processing accelerator core operations being specified by said corresponding memory address location.

2. The computer system of claim 1, **characterized in that** said translator (34) is implemented by translation instructions executed by said native processing unit (16).

3. The computer system of claim 1 or 2, **characterized in that** said core control logic (14, 40, 46) includes:

a stack controller (40) for controlling said writing and reading of data to and from said stack cache memory device (38) in accordance with a stack methodology; and

a decoder (46) for decoding each of said selected address values (60, 62, 64) for the purpose of determining said corresponding processing accelerator core operation.

4. The computer system of any of claims 1 - 3, **characterized by** a main memory unit (30), and wherein said core control logic (14, 40, 46) further comprises a direct memory access controller (14) operative to shift data from said stack cache memory device (38) to said main memory unit (30) when said stack cache memory device (38) ap-

proaches a stack overflow condition, and also operative to load data from said main memory unit (30) to said stack cache memory device (38) when said stack cache memory device (38) approaches a stack underflow condition.

- 5 5. The computer system of claim 4, **characterized in that** said direct memory access controller (14) is further operative to:

load a local variable from said main memory unit (30) to said stack cache memory device (38) in the event of a Variable Load Miss;
 store a local variable to said main memory unit (30) in the event of a Variable Store miss;
 10 transfer data from said stack cache memory device (38) to said main memory unit (30) in order to prepare for a context switch operation; and
 perform accelerating thread context switching by moving data into and out of said stack cache memory device (38) as required.

- 15 6. The computer system of any of claims 1 - 5, **characterized in that** each of said selected address values (60, 62, 64) includes:

20 a plurality of select area bits (62) for selecting from a plurality of memory spaces areas (52, 54, 56, 58) of said stack cache memory device (38), each of said memory space areas (52, 54, 56, 58) including a particular set of said memory address locations; and
 a plurality of select bits (64) for selecting particular ones of each of said memory address locations of each of said memory space areas (52, 54, 56, 58), and for determining the operation to be performed.

- 25 7. The computer system of any of claims 1 - 6, **characterized in that** said interpretive language is Java™.

8. The computer system of any of claims 1 - 7, **characterized in that** said processing accelerator core operations include arithmetic operations.

- 30 9. The computer system of any of claims 1 - 8, **characterized in that** said processing accelerator core operations include Boolean logic operations.

10. A method of providing accelerated processing of stack oriented interpretive language instructions in a computer system (20) including a native processing unit (16) for executing native language instructions,

35 **characterized by** said computer system (20) comprising a processing accelerator core (10) coupled for communication with the native processing unit (16), the processing accelerator core (10) including a stack cache memory device (38) having a plurality of memory address locations, control logic (14, 40, 46) operative to control writing and reading of data to and from the stack cache memory device (38), and an arithmetic logic unit (36) for receiving data from corresponding ones of said memory address locations of said stack cache device (38), and being operative to perform processing accelerator core operations using said data, and said method comprising the steps of:

40 receiving interpretive language instructions;
 translating (76) said interpretive language instructions to native language instructions;
 45 generating selected address values (60, 62, 64) associated with selected ones of said native language instructions, each of said selected instructions and said corresponding selected address values (60, 62, 64) being associated with a corresponding processing accelerator core operation;
 writing and reading data associated with said selected native language instructions to and from locations of the stack cache memory device (38), said memory address locations being indicated by corresponding ones of said selected address values (60, 62, 64), said data being associated with corresponding ones of said processing accelerator core operations; and
 providing said data from corresponding ones of the memory address locations of the stack cache memory device (38) to the arithmetic logic unit (36), and using the arithmetic logic unit (36) to perform said processing accelerator core operations using said data, each of said processing accelerator core operations being specified by said corresponding memory address location.

- 50 11. The method of claim 10, **characterized in that** said step of translating (76) includes executing instructions using the native processing unit (16).

12. The method of claim 10 or 11, **characterized in that** said core control logic (14, 40, 46) includes:

a stack controller (40) for controlling said writing and reading of data to and from said stack cache memory device (38) in accordance with a stack methodology; and
 5 a decoder (46) for decoding each of said selected address values (60, 62, 64) for the purpose of determining said corresponding core operation.

13. The method of any of claims 10 - 12, **characterized by** direct memory access control steps of:

shifting data from the stack cache memory device (38) to a main memory unit (30) of the computer system (20) when the stack cache memory device (38) approaches a stack overflow condition; and
 10 loading data from said main memory unit (30) to said stack cache memory device (38) when the stack cache memory device (38) approaches a stack underflow condition.

15 14. The method of claim 13, **characterized by** direct memory access control steps of:

loading a local variable from said main memory unit (30) to said stack cache memory device (38) in the event of a Variable Load Miss;
 20 storing a local variable to said main memory unit (30) in the event of a Variable Store miss;
 transferring data from said stack cache memory device (38) to said main memory unit (30) in order to prepare for a context switch operation; and
 performing accelerating thread context switching by moving data into and out of said stack cache memory device (38) as required.

25 15. The method of any of claims 10 - 14, **characterized in that** each of said selected address values (60, 62, 64) includes:

a plurality of select area bits (62) for selecting from a plurality of memory spaces areas (52, 54, 56, 58) of said stack cache memory device (38), each of said memory space areas (52, 54, 56, 58) including a particular set of said memory address locations; and
 30 a plurality of select bits (64) for selecting particular ones of each of said memory address locations of each of said memory space areas (52, 54, 56, 58), and for determining the processing accelerator core operation to be performed.

35 16. The method of any of claims 10 - 15, **characterized in that** said interpretive language is Java".

17. The method of any of claims 10 - 16, **characterized in that** said processing accelerator core operations include arithmetic operations.

40 18. The method of any of claims 10 - 17, **characterized in that** said processing accelerator core operations include Boolean logic operations.

Patentansprüche

45 1. Computersystem, das eine beschleunigte Verarbeitung von stapelorientierten interpretativen Sprachbefehlen bereitstellt, wobei das System (20) aufweist:

eine native Verarbeitungseinheit (16) zum Ausführen von nativen Sprachbefehlen;
 50 einen Übersetzer (34) zum Übersetzen der interpretativen Sprachbefehle in native Sprachbefehle;
gekennzeichnet durch
 die Tatsache, dass der Übersetzer (34) dazu in der Lage ist, ausgewählte Adresswerte (60, 62, 64) zu erzeugen, die ausgewählten Befehlen der nativen Sprachbefehle zugeordnet sind, wobei jeder der ausgewählten Befehle und der entsprechenden ausgewählten Adresswerte (60, 62, 64) einer entsprechenden Verarbeitungsbeschleunigungskern-Operation zugeordnet sind;
 55 einem Verarbeitungsbeschleunigungskern (10), der zur Kommunikation mit der nativen Verarbeitungseinheit (16) angeschlossen ist, wobei der Verarbeitungsbeschleunigungskern (10) aufweist
 eine Cache-Stapelspeichereinrichtung (38), die eine Vielzahl von Speicheradresssorten aufweist, die

durch entsprechende Werte der ausgewählten Adresswerte (60, 62, 64) indiziert sind, wobei die Speicheradressorte zum Speichern von Daten dienen, und zwar in Zuordnung zu entsprechenden Operationen der Verarbeitungsbeschleunigungskern-Operationen,

5 eine Kernsteuerlogik (14, 40, 46), die auf die ausgewählten Adresswerte (60, 62, 64) anspricht und betriebsbereit ist, um das Schreiben und das Lesen der Daten, die den ausgewählten nativen Sprachbefehlen zugeordnet sind, in die Cache-Stapelspeichereinrichtung (38) hinein und aus dieser heraus zu steuern, und

10 eine Arithmetiklogikeinheit (36) zum Empfangen der Daten von entsprechenden Orten der Speicheradressorte der Cache-Stapelspeichereinrichtung (38), wobei die Arithmetiklogikeinheit (36) betriebsbereit ist, die Verarbeitungsbeschleunigungskern-Operationen unter Verwendung der Daten durchzuführen, wobei jede der Verarbeitungsbeschleunigungskern-Operationen durch den entsprechenden Speicheradressort spezifiziert ist.

2. Computersystem nach Anspruch 1, dadurch gekennzeichnet, dass der Übersetzer (34) durch Übersetzungsbefehle implementiert ist, die von der nativen Verarbeitungseinheit (16) ausgeführt werden.

3. Computersystem nach Anspruch 1 oder 2, dadurch gekennzeichnet, dass die Kernsteuerlogik (14, 40, 46) aufweist:

20 eine Stapelsteuereinrichtung (40) zum Steuern des Schreibens und Lesens von Daten in die Cache-Stapelspeichereinrichtung (38) bzw. aus dieser heraus, und zwar gemäß einer Stapelmethodologie; und einen Decodierer (46) zum Decodieren von jedem der ausgewählten Adresswerte (60, 62, 64) zum Zwecke des Bestimmens der entsprechenden Verarbeitungsbeschleunigungskern-Operation.

4. Computersystem nach einem der Ansprüche 1 bis 3, gekennzeichnet durch eine Hauptspeichereinheit (30), und wobei die Kernsteuerlogik (14, 40, 46) ferner eine Direktspeicherzugriff-Steuereinrichtung (14) aufweist, die betriebsbereit ist, Daten von der Cache-Stapelspeichereinrichtung (38) in die Hauptspeichereinheit (30) zu verschieben, wenn die Cache-Stapelspeichereinrichtung (38) sich einem Stapelüberlaufzustand nähert, und die ferner betriebsbereit ist, Daten aus der Hauptspeichereinheit (30) in die Cache-Stapelspeichereinrichtung (38) zu laden, wenn die Cache-Stapelspeichereinrichtung (38) sich einem Stapelunterlaufzustand nähert.

5. Computersystem nach Anspruch 4, dadurch gekennzeichnet, dass die Direktspeicherzugriff-Steuereinrichtung (14) ferner betriebsbereit ist:

35 eine lokale Variable aus der Hauptspeichereinheit (30) in die Cache-Stapelspeichereinrichtung (38) zu laden, und zwar für den Fall eines Variablen-Ladefehlschlages ("Variable Load Miss");

eine lokale Variable in die Hauptspeichereinheit (30) zu speichern, und zwar für den Fall eines Variablen-Speicherfehlschlages ("Variable Store Miss");

Daten von der Cache-Stapelspeichereinrichtung (38) in die Hauptspeichereinheit (30) zu übertragen, um Vorbereitungen für eine Kontext-Umschaltoperation zu treffen; und

40 ein Verkettungskontext-Umschalten zu beschleunigen, indem nach Erfordernis Daten in die Cache-Stapelspeichereinrichtung (38) hinein bzw. aus dieser heraus verschoben werden.

6. Computersystem nach einem der Ansprüche 1 bis 5, dadurch gekennzeichnet, dass jeder der ausgewählten Adresswerte (60, 62, 64) aufweist:

45 eine Vielzahl von Bereichsauswahlbits (62) zum Auswählen aus einer Vielzahl von Speicherbereichen (52, 54, 56, 58) der Cache-Stapelspeichereinrichtung (38), wobei jeder der Speicherbereiche (52, 54, 56, 58) einen bestimmten Satz der Speicheradressorte beinhaltet; und

50 eine Vielzahl von Auswahlbits (64) zum Auswählen von bestimmten Adressorten der Speicheradressorte von jedem der Speicherbereiche (52, 54, 56, 58), und zum Bestimmen der durchzuführenden Operation.

7. Computersystem nach einem der Ansprüche 1 bis 6, dadurch gekennzeichnet, dass die interpretative Sprache Java™ ist.

8. Computersystem nach einem der Ansprüche 1 bis 7, dadurch gekennzeichnet, dass die Verarbeitungsbeschleunigungskern-Operationen arithmetische Operationen aufweisen.

9. Computersystem nach einem der Ansprüche 1 bis 8, dadurch gekennzeichnet, dass die Verarbeitungsbeschleu-

nigungskern-Operationen Bool'sche Logikoperationen aufweisen.

10. Verfahren zum Bereitstellen einer beschleunigten Verarbeitung von stapelorientierten interpretativen Sprachbefehlen in einem Computersystem (20), das eine native Verarbeitungseinheit (16) zum Ausführen von nativen Sprachbefehlen aufweist,

gekennzeichnet durch

die Tatsache, dass das Computersystem (20) einen Verarbeitungsbeschleunigungskern (10) aufweist, der zur Kommunikation mit der nativen Verarbeitungseinheit (16) angeschlossen ist, wobei der Verarbeitungsbeschleunigungskern (10) eine Cache-Stapelspeichereinrichtung (38) aufweist, die eine Vielzahl von Speicheradressorten beinhaltet, eine Steuerlogik (14, 40, 46) aufweist, die betriebsbereit ist, das Schreiben in und das Lesen aus der Cache-Stapelspeichereinrichtung (38) zu steuern, und eine Arithmetiklogikeinheit (36) aufweist, zum Empfangen von Daten aus entsprechenden Orten der Speicheradressorte der Cache-Stapelspeichereinrichtung (38), und die betriebsbereit ist, unter Verwendung der Daten Verarbeitungsbeschleunigungskern-Operationen durchzuführen, und wobei das Verfahren die Schritte aufweist:

15. Empfangen von interpretativen Sprachbefehlen;

Übersetzen (76) der interpretativen Sprachbefehle in native Sprachbefehle;

Erzeugen von ausgewählten Adresswerten (60, 62, 64), die ausgewählten Befehlen der nativen Sprachbefehle zugeordnet sind, wobei die ausgewählten Befehle und die entsprechenden ausgewählten Adresswerte (60, 62, 64) jeweils einer entsprechenden Verarbeitungsbeschleunigungskern-Operation zugeordnet sind;

Schreiben und Lesen von Daten, die den ausgewählten nativen Sprachbefehlen zugeordnet sind, in Orte der Cache-Stapelspeichereinrichtung (38) bzw. aus dieser heraus, wobei die Speicheradressorte **durch** entsprechende Werte der ausgewählten Adresswerte (60, 62, 64) indiziert sind, wobei die Daten zugeordnet sind zu entsprechenden Operationen der Verarbeitungsbeschleunigungskern-Operationen; und

Bereitstellen der Daten aus entsprechenden Orten der Speicheradressorte der Cache-Stapelspeichereinrichtung (38) an die Arithmetiklogikeinheit (36), und Verwenden der Arithmetiklogikeinheit (36) dazu, die Verarbeitungsbeschleunigungskern-Operationen unter Verwendung der Daten durchzuführen, wobei jede der Verarbeitungsbeschleunigungskern-Operationen **durch** den entsprechenden Speicheradressort spezifiziert ist.

30. 11. Verfahren nach Anspruch 10, **dadurch gekennzeichnet, dass** der Schritt des Übersetzens (76) das Ausführen von Befehlen unter Verwendung der nativen Verarbeitungseinheit (16) beinhaltet.

12. Verfahren nach Anspruch 10 oder 11, **dadurch gekennzeichnet, dass** die Kernsteuerlogik (14, 40, 46) aufweist:

35. eine Stapelsteuereinrichtung (40) zum Steuern des Schreibens und Lesens von Daten in die Cache-Stapelspeichereinrichtung (38) bzw. aus dieser heraus, und zwar gemäß einer Stapel-Methodologie; und einen Decodierer (46) zum Decodieren von jedem der ausgewählten Adresswerte (60, 62, 64) zum Zwecke des Bestimmens der entsprechenden Kern-Operation.

40. 13. Verfahren nach einem der Ansprüche 10 bis 12, **gekennzeichnet durch** Direktspeicherzugriffs-Steuerschritte:

Verschieben von Daten aus der Cache-Stapelspeichereinrichtung (38) in eine Hauptspeichereinheit (30) des Computersystems (20), wenn sich die Cache-Stapelspeichereinrichtung (38) einem Stapelüberlaufzustand nähert; und

45. Laden von Daten aus der Hauptspeichereinheit (30) in die Cache-Stapelspeichereinrichtung (38), wenn sich die Cache-Stapelspeichereinrichtung (38) einem Stapelunterlaufzustand nähert.

14. Verfahren nach Anspruch 13, **gekennzeichnet durch** Direktspeicherzugriffs-Steuerschritte:

50. Laden einer lokalen Variable aus der Hauptspeichereinheit (30) in die Cache-Stapelspeichereinrichtung (38) für den Fall eines Variablen-Ladefehlschlages;

Speichern einer lokalen Variable in die Hauptspeichereinheit (30) für den Fall eines Variablen-Speicherfehlschlages;

Übertragen von Daten aus der Cache-Stapelspeichereinrichtung (38) in die Hauptspeichereinheit (30), um Vorbereitungen für eine Kontext-Umschaltoperation zu treffen; und

55. Durchführen von beschleunigtem Verkettungskontext-Umschalten **durch** Verschieben von Daten in die Cache-Stapelspeichereinrichtung (38) bzw. aus dieser heraus, je nach Anforderung.

15. Verfahren nach einem der Ansprüche 10 bis 14, **dadurch gekennzeichnet, dass** jeder der ausgewählten Adresswerte (60, 62, 64) aufweist:

5 eine Vielzahl von Bereichsauswahlbits (62) zum Auswählen aus einer Vielzahl von Speicherbereichen (52, 54, 56, 58) der Cache-Stapelspeichereinrichtung (38), wobei jeder der Speicherbereiche (52, 54, 56, 58) einen bestimmten Satz der Speicheradressorte beinhaltet; und
10 eine Vielzahl von Auswahlbits (64) zum Auswählen von bestimmten Orten der Speicheradressorte von jedem der Speicherbereiche (52, 54, 56, 58), und zum Bestimmen der durchzuführenden Verarbeitungsbeschleunigungskern-Operation.

- 10 16. Verfahren nach einem der Ansprüche 10 bis 15, **dadurch gekennzeichnet, dass** die interpretative Sprache Java™ ist.

- 15 17. Verfahren nach einem der Ansprüche 10 bis 16, **dadurch gekennzeichnet, dass** die Verarbeitungsbeschleunigungskern-Operationen arithmetische Operationen beinhalten.

- 20 18. Verfahren nach einem der Ansprüche 10 bis 17, **dadurch gekennzeichnet, dass** die Verarbeitungsbeschleunigungskern-Operationen Bool'sche Logikoperationen beinhalten.

Revendications

1. Système informatique fournissant un traitement accéléré d'instructions de langage interprétatif orientées en pile, ledit système (20) comprenant :

25 une unité de traitement natif (16) pour exécuter des instructions en langage natif ;
 un traducteur (34) pour traduire les instructions en langage interprétatif en instructions en langage natif ;

30 **caractérisé par le fait que** ledit traducteur (34) est capable de générer des valeurs d'adresse sélectionnées (60, 62, 64) associées à celles sélectionnées desdites instructions en langage natif, chacune desdites instructions sélectionnées et desdites valeurs d'adresse sélectionnées (60, 62, 64) correspondantes étant associée à une opération correspondante d'un cœur d'accélérateur de traitement ;

35 un cœur d'accélérateur de traitement (10) couplé pour communiquer avec ladite unité de traitement natif (16), ledit cœur d'accélérateur de traitement (10) comprenant

 un dispositif (38) à mémoire cache de pile ayant une pluralité de positions d'adresse mémoire indiquées par celles correspondantes desdites valeurs d'adresse sélectionnées (60, 62, 64), lesdites positions d'adresse mémoire étant destinées à stocker des données associées à celles correspondantes desdites opérations du cœur d'accélérateur de traitement,

40 une logique de contrôle de cœur (14, 40, 46) réagissant auxdites valeurs d'adresse sélectionnées (60, 62, 64) et fonctionnant pour contrôler l'écriture et la lecture desdites données associées auxdites instructions en langage natif sélectionnées vers et depuis ledit dispositif (38) à mémoire cache de pile, et

45 une unité logique arithmétique (36) pour recevoir lesdites données depuis celles correspondantes desdites positions d'adresse mémoire dudit dispositif (38) à mémoire cache de pile, et fonctionnant pour exécuter lesdites opérations de cœur d'accélérateur de traitement en utilisant lesdites données, chacune desdites opérations de cœur d'accélérateur de traitement étant spécifiée par ladite position d'adresse mémoire correspondante.

2. Système informatique selon la revendication 1, **caractérisé en ce que** ledit traducteur (34) est réalisé par des instructions de traduction exécutées par ladite unité de traitement natif (16).

- 50 3. Système informatique selon la revendication 1 ou 2, **caractérisé en ce que** la logique de contrôle de cœur (14, 40, 46) comprend :

55 un contrôleur de pile (40) pour contrôler ladite écriture et lecture vers et depuis ledit dispositif (38) à mémoire cache de pile conformément à la méthodologie de pile ; et

 un décodeur (46) pour décoder chacune desdites valeurs d'adresse sélectionnées (60, 62, 64) afin de déterminer ladite opération de cœur d'accélérateur de traitement correspondante.

4. Système informatique selon l'une quelconque des revendications 1 à 3, **caractérisé par** une unité de mémoire

principale (30), et dans lequel ladite logique de contrôle de cœur (14, 40, 46) comprend de plus un contrôleur d'accès direct en mémoire (14) servant à décaler des données du dispositif (38) à mémoire cache de pile vers ladite unité de mémoire principale (30) quand ledit dispositif (38) à mémoire cache de pile approche d'un état de dépassement de pile, et fonctionnant également pour charger des données de ladite unité de mémoire principale (30) vers ledit dispositif (38) à mémoire cache de pile quand ledit dispositif (38) à mémoire cache de pile approche d'un état de souspassement de pile.

- 5 5. Système informatique selon la revendication 4, **caractérisé en ce que** ledit contrôleur d'accès direct en mémoire (14) opère de plus pour :

10 charger une variable locale de ladite unité de mémoire principale (30) vers ledit dispositif (38) à mémoire cache de pile en cas d'erreur de chargement de variable ;
 charger une variable locale dans ladite unité de mémoire principale (30) en cas d'erreur de stockage de variable ;
 15 transférer des données du dispositif (38) à mémoire cache de pile à la dite unité de mémoire principale (30) afin de préparer une opération de commutation de contexte ; et
 effectuer une commutation de contexte de thread accélérée en déplaçant des données dans et depuis ledit dispositif (38) à mémoire cache de pile, selon les besoins.

- 20 6. Système informatique selon l'une quelconque des revendications 1 à 5, **caractérisé en ce que** chacune desdites valeurs d'adresse sélectionnées (60, 62, 64) comprend :

25 une pluralité de bits de zone choisie (62) pour sélectionner parmi une pluralité de zones d'espace mémoire (52, 54, 56, 58) dudit dispositif (38) à mémoire cache de pile, chacune desdites zones d'espace mémoire (52, 54, 56, 58) comprenant un ensemble particulier desdites positions d'adresse mémoire ; et
 une pluralité de bits de sélection (64) pour sélectionner celles particulières desdites positions d'adresse mémoire de chacune desdites zones d'espace mémoire (52, 54, 56, 58) et pour déterminer l'opération à effectuer.

- 30 7. Système informatique selon l'une quelconque des revendications 1 à 6, **caractérisé en ce que** le langage interpréitatif est Java™.

- 35 8. Système informatique selon l'une quelconque des revendications 1 à 7, **caractérisé en ce que** lesdites opérations de cœur d'accélérateur de traitement comprennent des opérations arithmétiques.

- 35 9. Système informatique selon l'une quelconque des revendications 1 à 8, **caractérisé en ce que** lesdites opérations de cœur d'accélérateur de traitement comprennent des opérations logique booléennes.

- 40 10. Procédé pour fournir un traitement accéléré d'instructions de langage interpréitatif orienté pile dans un système informatique (20) comprenant une unité de traitement natif (16) pour exécuter des instructions en langage natif, **caractérisé par le fait que** ledit système informatique (20) comprend un cœur d'accélérateur de traitement (10) couplé en communication avec l'unité de traitement natif (16), le cœur d'accélérateur de traitement (10) comprenant un dispositif (38) à mémoire cache de pile ayant une pluralité de positions d'adresse mémoire, une logique de contrôle (14, 40, 46) servant à contrôler l'écriture et la lecture de données vers et depuis le dispositif (38) à mémoire cache de pile, et une unité logique arithmétique (36) pour recevoir des données depuis celles correspondantes desdites positions d'adresse mémoire dudit dispositif (38) à mémoire cache de pile, et fonctionnant pour exécuter des opérations de cœur d'accélérateur de traitement à l'aide desdites données, et ledit procédé comprenant les étapes consistant à :

50 recevoir des instructions en langage interpréitatif ;
 traduire (76) lesdites instructions en langage interpréitatif en des instructions en langage natif ;
 générer des valeurs d'adresse sélectionnées (60, 62, 64) associées à celles sélectionnées desdites instructions en langage natif, chacune desdites instructions sélectionnées et desdites valeurs d'adresse mémoire (60, 62, 64) étant associée à une opération de cœur d'accélérateur de traitement correspondante ;
 écrire et lire des données associées auxdites instructions en langage natif sélectionnées vers et depuis des positions du dispositif (38) à mémoire cache de pile, lesdites positions d'adresse mémoire étant indiquées par celles correspondantes des valeurs d'adresse sélectionnées (60, 62, 64), lesdites données étant associées à celles correspondantes desdites opérations de cœur d'accélérateur de traitement ; et
 fournir lesdites données depuis celles correspondantes des positions d'adresse mémoire du dispositif (38) à

5 mémoire cache de pile à l'unité logique arithmétique (36), et utiliser l'unité logique arithmétique (36) pour effectuer lesdites opérations de coeur d'accélérateur de traitement en utilisant lesdites données, chacune desdites opérations de coeur d'accélérateur de traitement étant spécifiée par ladite position d'adresse mémoire correspondante.

10 11. Procédé selon la revendication 10, **caractérisé en ce que** ladite étape de traduction (76) comprend l'exécution d'instructions à l'aide de l'unité de traitement natif (16).

15 12. Procédé selon la revendication 10 ou 11, **caractérisé en ce que** ladite logique de contrôle de coeur(14, 40, 46) comprend :

un contrôleur de pile (40) pour contrôler ladite écriture et lecture de données vers et depuis ledit dispositif (38) à mémoire cache de pile conformément à une méthodologie de pile ; et
16 un décodeur (46) pour décoder chacune desdites valeurs d'adresse sélectionnées (60, 62, 64) afin de déterminer ladite opération de coeur correspondante.

20 13. Procédé selon l'une des revendications 10 à 12, **caractérisé par** des étapes de contrôle d'accès direct en mémoire consistant à :

décaler des données du dispositif (38) à mémoire cache de pile vers une unité de mémoire principale (30) du système informatique (20) quand le dispositif (38) à mémoire cache de pile approche d'un état de dépassement de pile ; et
25 charger des données de ladite unité de mémoire principale (30) vers ledit dispositif (38) à mémoire cache de pile quand ledit dispositif (38) à mémoire cache de pile approche d'un état de souppassement de pile.

25 14. Procédé selon la revendication 13, **caractérisé par** des étapes de contrôle d'accès direct en mémoire consistant à :

charger une variable locale de ladite unité de mémoire principale (30) vers ledit dispositif (38) à mémoire cache de pile en cas d'erreur de chargement de variable ;
30 charger une variable locale dans ladite unité de mémoire principale (30) en cas d'erreur de stockage de variable ;
transférer des données du dispositif (38) à mémoire cache de pile à la dite unité de mémoire principale (30) afin de préparer une opération de commutation de contexte ; et
35 effectuer une commutation de contexte de fil de discussion accélérée en déplaçant des données dans et depuis ledit dispositif (38) à mémoire cache de pile, selon les besoins.

40 15. Procédé selon l'une quelconque des revendications 10 à 14, **caractérisé en ce que** chacune desdites valeurs d'adresse sélectionnées (60, 62, 64) comprend :

une pluralité de bits de zone choisie (62) pour sélectionner parmi une pluralité de zones d'espace mémoire (52, 54, 56, 58) dudit dispositif (38) à mémoire cache de pile, chacune desdites zones d'espace mémoire (52, 54, 56, 58) comprenant un ensemble particulier desdites positions d'adresse mémoire ; et
45 une pluralité de bits de sélection (64) pour sélectionner celles particulières desdites positions d'adresse mémoire de chacune desdites zones d'espace mémoire (52, 54, 56, 58) et pour déterminer l'opération de coeur d'accélérateur de traitement à effectuer.

46 16. Procédé selon l'une quelconque des revendications 10 à 15, **caractérisé en ce que** ledit langage interprétatif est Java™.

50 17. Procédé selon l'une quelconque des revendications 10 à 16, **caractérisé en ce que** lesdites opérations de coeur d'accélérateur de traitement comprennent des opérations arithmétiques.

55 18. Procédé selon l'une quelconque des revendications 10 à 17, **caractérisé en ce que** lesdites opérations de coeur d'accélérateur de traitement comprennent des opérations logique booléennes.

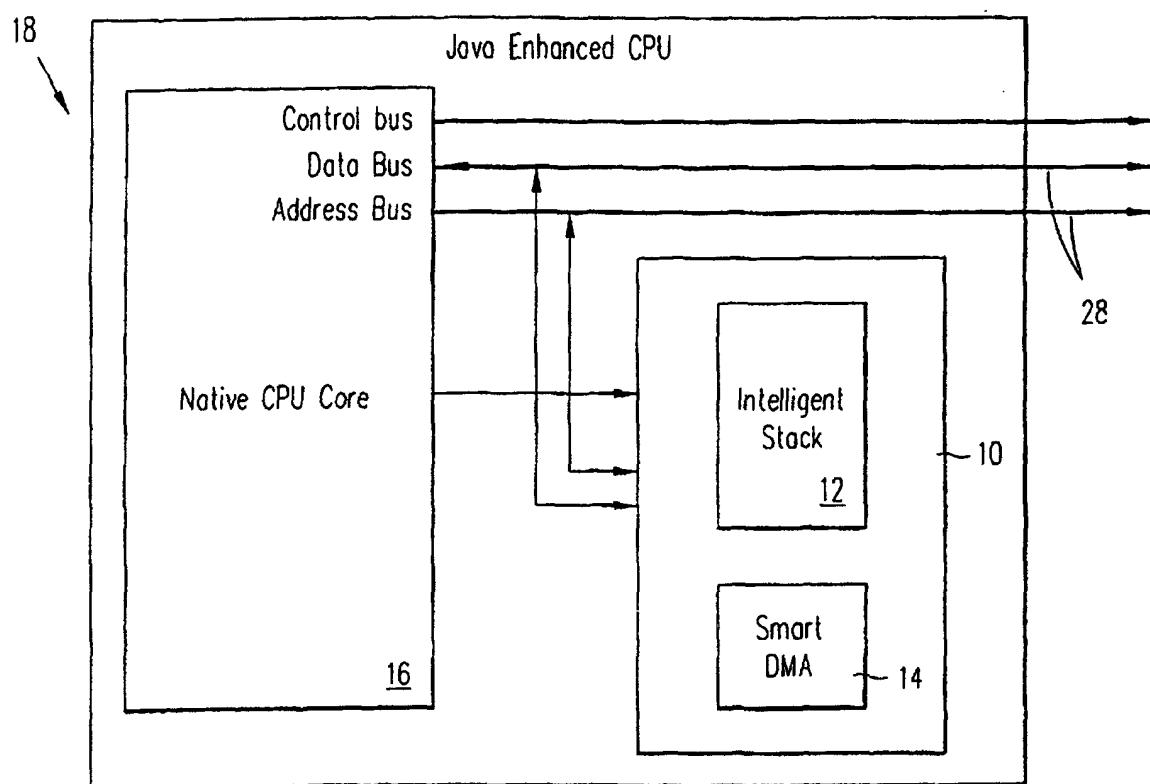


FIG. 1

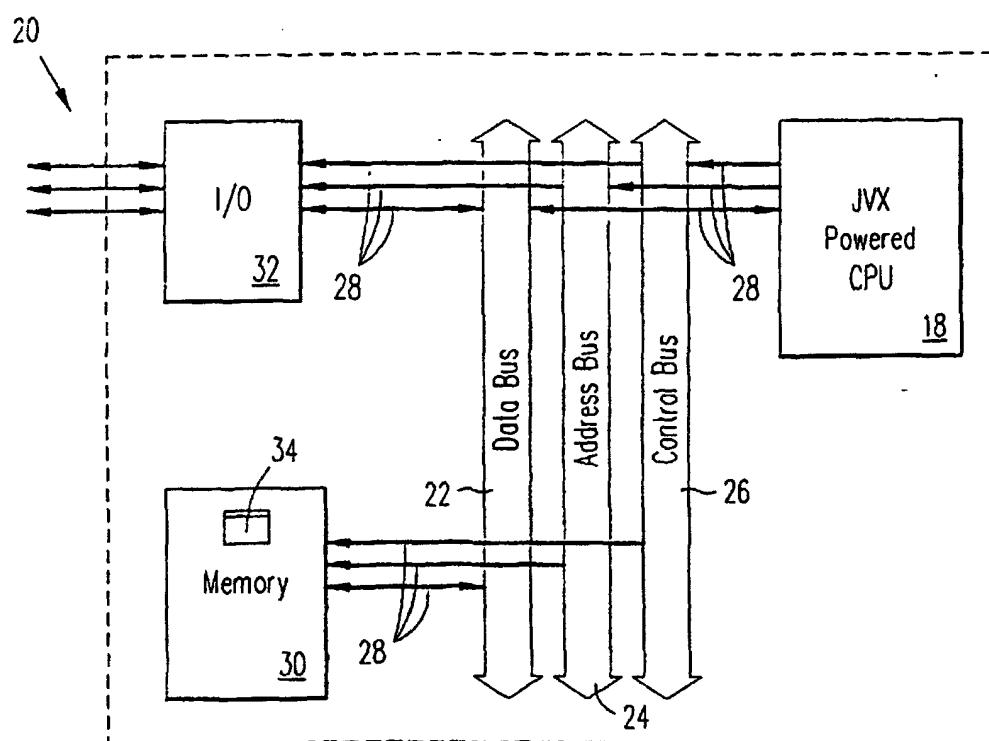


FIG. 2

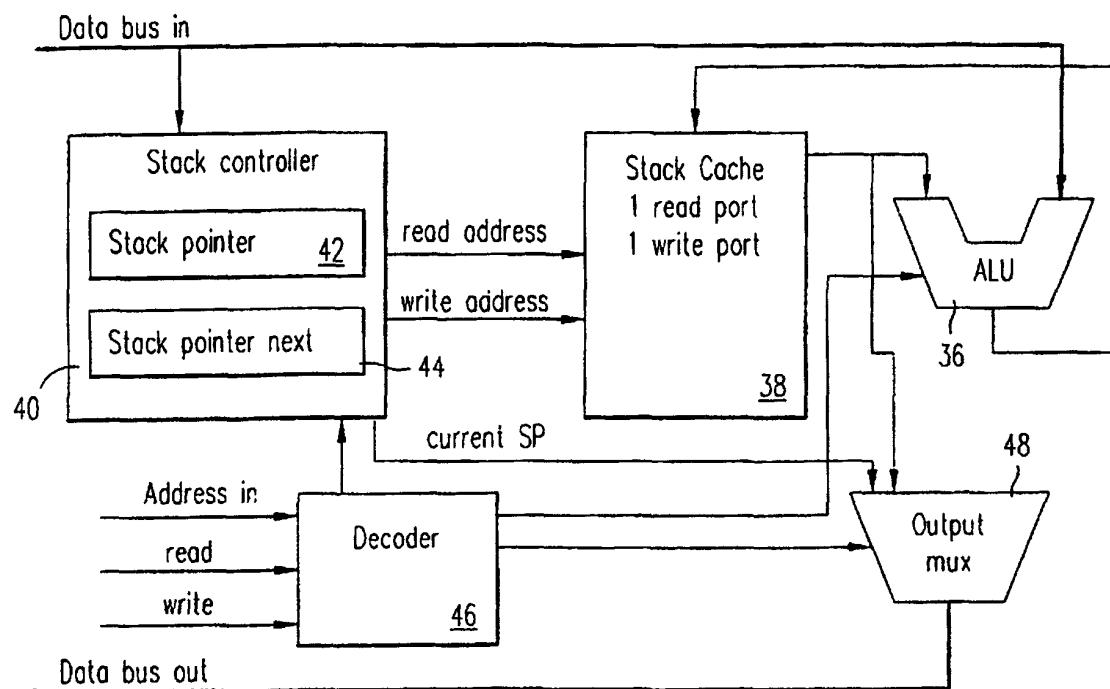


FIG. 3

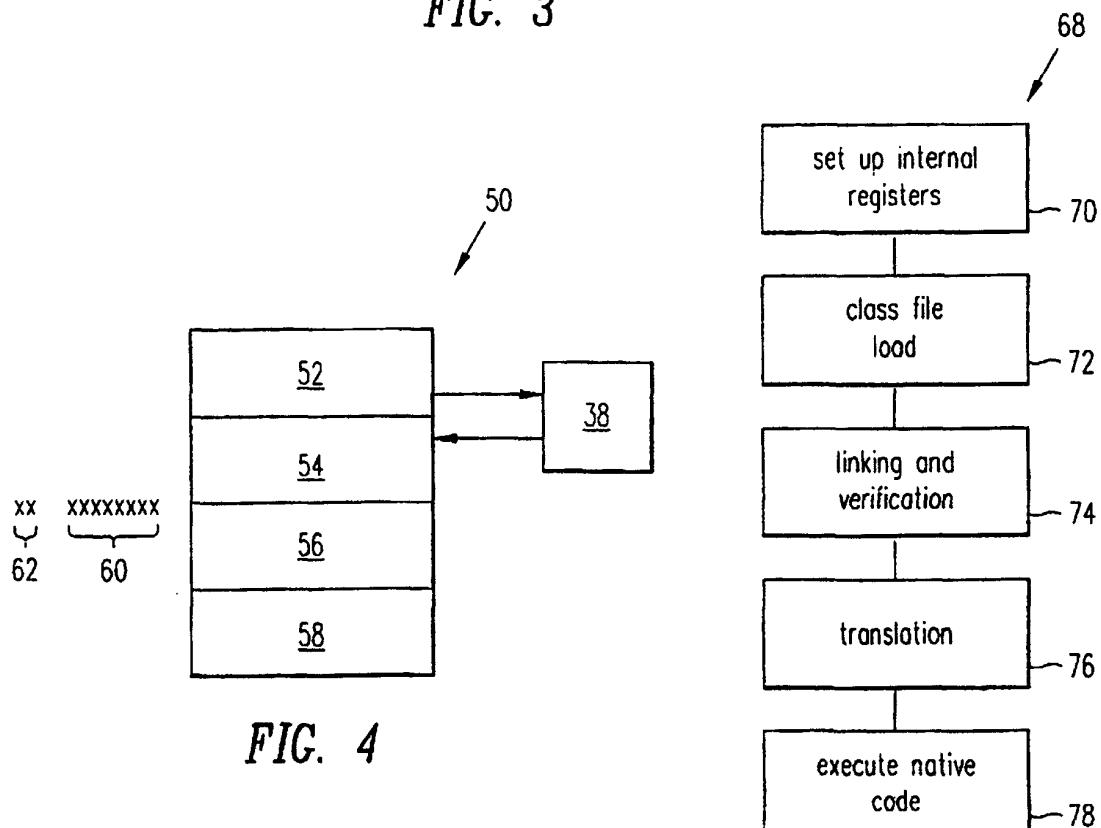


FIG. 4

FIG. 5

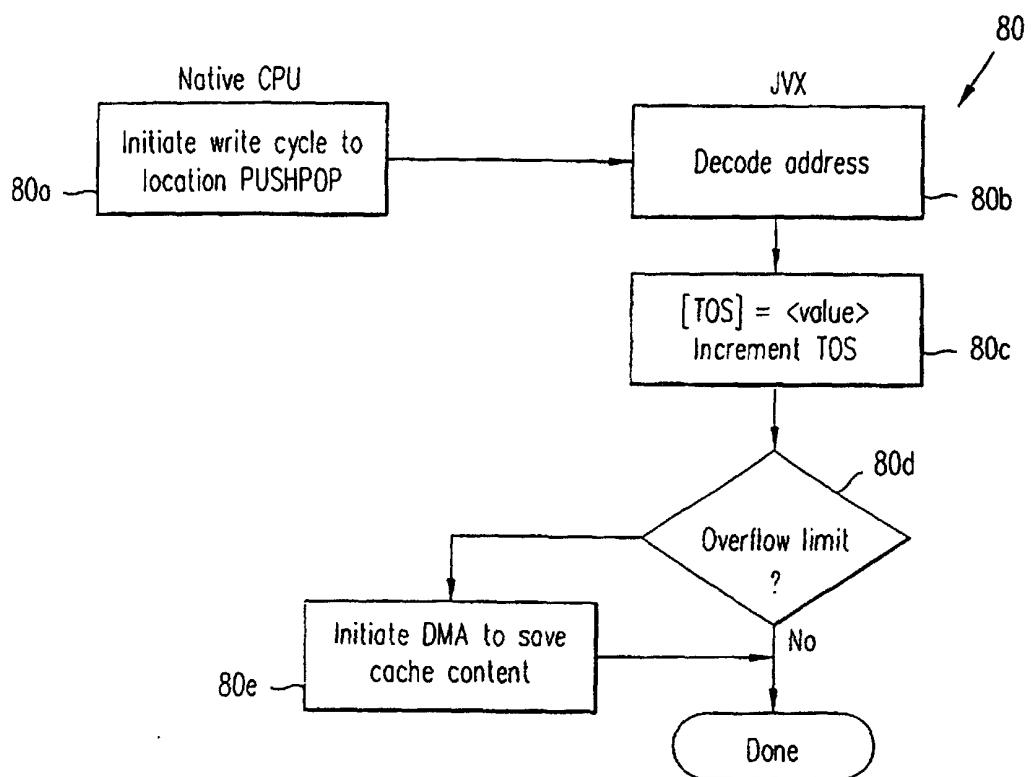


FIG. 6

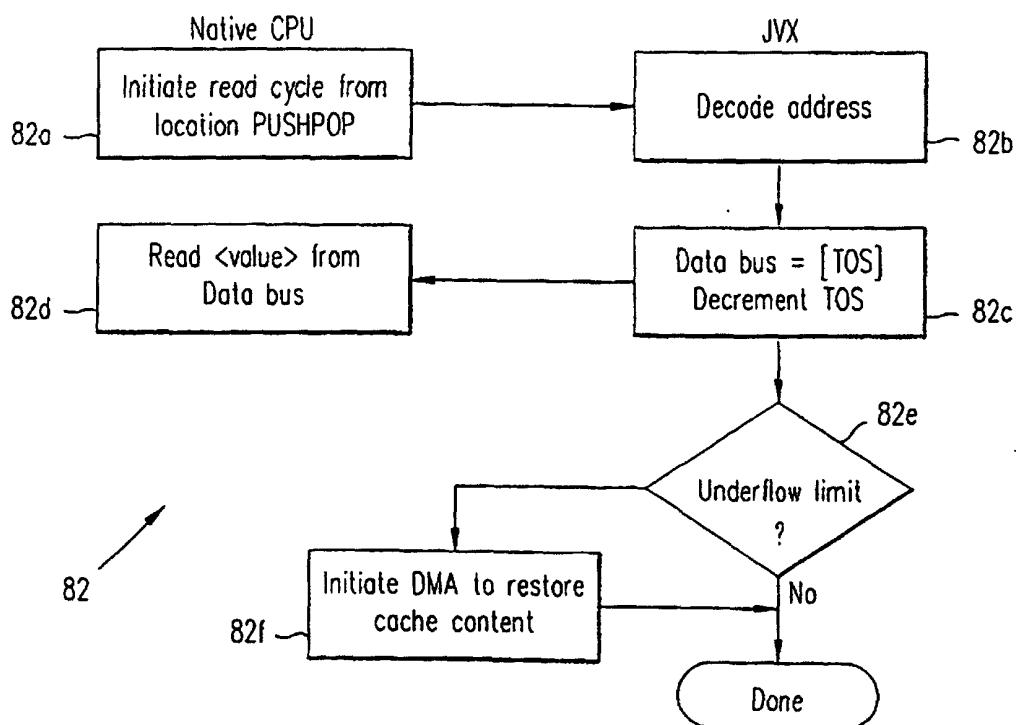


FIG. 7

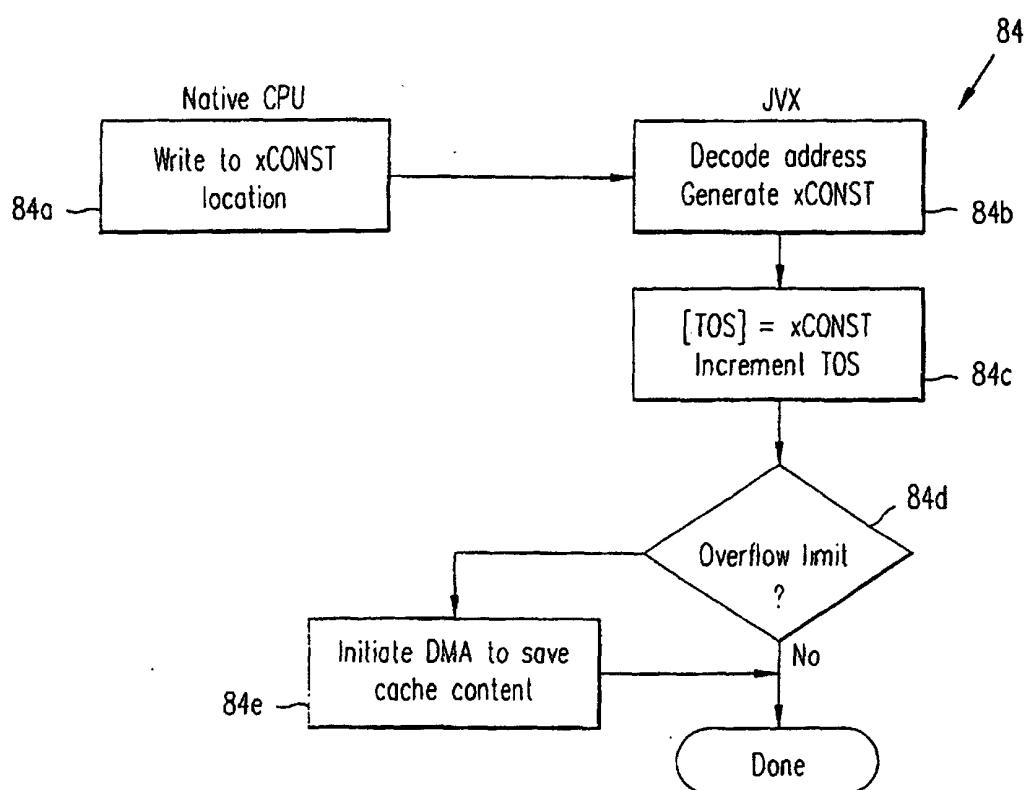


FIG. 8

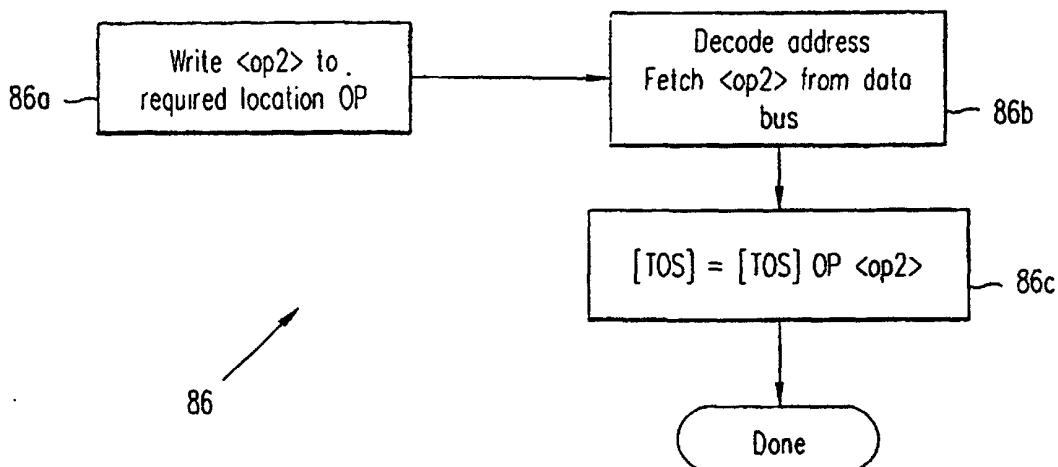


FIG. 9

